PCT/AU2004/001053

AUSTRALIA

REC'D 24 AUG 2004

Patent Office
Canberra

PCT

I, JULIE BILLINGSLEY, TEAM LEADER EXAMINATION SUPPORT AND SALES hereby certify that annexed is a true copy of the Provisional specification in connection with Application No. 2003904170 for a patent by CLIPSAL INTEGRATED SYSTEMS PTY LTD as filed on 08 August 2003.

COMMONWEALTH OF AUSTRALIA
AUSTRALIA
PATENT OFFICE

WITNESS my hand this
Seventeenth day of August 2004

J. Billingsley

JULIE BILLINGSLEY
TEAM LEADER EXAMINATION
SUPPORT AND SALES

PRIORITY DOCUMENT
SUBMITTED OR TRANSMITTED IN COMPLIANCE WITH RULE 17.1(a) OR (b)

CLIPSAL INTEGRATED SYSTEMS PTY LTD

AUSTRALIA
PATENTS ACT 1990

PROVISIONAL SPECIFICATION FOR THE INVENTION ENTITLED:

"RADIO NETWORK COMMUNICATION SYSTEM AND PROTOCOL"
This invention is described in the following statement:

TECHNICAL FIELD

This invention relates to a network of devices which communicate with each other via radio frequency.

BACKGROUND OF THE INVENTION

A network of devices can be created by arranging a group of devices which can communicate with each other via radio frequency (RF) means to transmit data between the devices. Provided that each of the devices is within the range of the maximum communications range of each device, each device can effectively communicate with each other device in the network.

In many such networks, the devices may all be transceivers, ie are each capable of both transmitting and receiving. For the purpose of transferring an item of data, one device acts predominantly as a transmitter while other devices act predominantly as receivers. In this context, transceivers, which act predominantly as transmitters will be referred to as transceiver/transmitters. Those which act predominantly as receivers will be referred to as transceiver/receivers.

The present invention is to be used in a "point to multipoint" system as opposed to a "point to point" system in which communications occur between only two devices at a time. In a point to multipoint communication system, communications occur between one device and two or more of the other devices in the network simultaneously.

A reliable "point to multipoint" communication system allows the creation of a shared network variable. This is a variable which is known to all of the devices in the network. For example, if one device wants to change the value of the shared network variable, it must transmit a request and be guaranteed that all devices receive and process the updated variable simultaneously. If the update cannot be made simultaneously, or not all other devices in the network receive the update, then the network does not have a shared network variable.

Shared network variables allow the creation of a network which has no central controller. All of the essential data about the operation and control of the network is known by each device in the network simultaneously. The data can be updated by any device in the network at any time and all other devices are guaranteed to update their data accordingly. This allows the control of devices within the network to be simplified, more flexible, and less costly when compared with networks having a central controller.

A single communication action between each of the devices is herein referred to as a Transaction. A Transaction occurs between a device (transceiver/transmitter) which transmits data to one or more transceiver/receivers of the data. The Transaction also includes data sent from the transceiver/receivers to the transceiver/transmitter as well as to each other of the transceiver/receiver devices in the network. This data return is generally an acknowledgement, indicating the success or failure of the receipt of the data sent by the transceiver/transmitter.
When transmitting to more than one transceiver/receiver simultaneously (also known as a broadcast or multicast), it is essential to know that all transceiver/receivers have successfully received the data. If even one transceiver/receiver has not successfully received the data (for example because of a bit error causing data corruption in one transceiver/receiver), then all other transceiver/receivers must be informed that not all of the other transceiver/receivers have successfully received the data. Each transceiver/receiver device must then discard the data it just received. This ensures that data used to update a shared network variable is received either by all transceiver/receivers, or by none or them.

Such networks can also optionally use a transmission system with dominant and inferior bits. This means that if there is a conflict, and two devices simultaneously transmit a dominant and inferior bit, then when monitoring the communication medium, each device will see the transmission of the dominant bit. The device transmitting the inferior bit knows that there has been a conflict and can take whatever action is appropriate. For example, this may mean the cessation of all further transmissions.

In conventional point to point communication protocols, a general procedure is to have each device transmit an acknowledge statement some time after receipt of the data. This has the drawback that the transmitting device must know exactly the number of receiving devices within the network, and how to contact each of them. The reliable transfer of the same piece of data to multiple receiving devices requires many transmissions of the same data, and a corresponding wait for each transmission to be acknowledged. The repetitive transmission of the same data to many recipients wastes the available bandwidth of the communication medium. This approach also requires the transmitter to obtain and store data about exactly which receiving devices are to accept a given transmission. This approach allows the creation of a shared network variable, at the expense of unnecessary complexity and poor use of the available bandwidth of the communication medium.

Alternatively a point to multipoint transmission can be used to transfer data to many recipients simultaneously, without any acknowledgment being returned. This renders the data transfer unreliable, and the transmitter will not be able to determine whether all of the receiving devices have successfully received the data. Unreliable transfer of data means that a shared network variable cannot be created.

It is an object of the present invention to provide a system and protocol for improving the communications between devices in an RF multicast communications system.

SUMMARY OF THE INVENTION

According to a first aspect of the present invention, there is provided a communications protocol for use in a network of devices, the protocol having a frame including a first time slot for transmitting data, a second time slot, after said first time slot, for transmitting a first acknowledgment state, and a third time slot, after said second time slot, for transmitting a second acknowledgement state.

One of the acknowledgment states is a positive acknowledge and the other acknowledgment state is a negative acknowledge.

Preferably, the first acknowledgment state is the positive acknowledge, and the second acknowledgment state is the negative acknowledge.

Preferably, the first time slot is variable in length and the second and third time slots are fixed in length. Preferably, the positive acknowledge transmission includes the transmission of a specific coded value containing sufficient redundancy to allow it to be recovered in the presence of received errors. Similarly the negative acknowledge transmission includes the transmission of a specific coded value containing sufficient redundancy to allow it to be recovered in the presence of received errors.

According to a second aspect of the present invention, there is provided a radio communications system including at least three transceivers, in use, one of which acts predominantly as a transmitter (transceiver/transmitter) while the other at least two transceivers act predominantly as receivers (transceiver/receiver), wherein the transceiver/transmitter transmits data in a first time slot to said transceiver/receivers, and wherein upon receipt of said data, each of said transceiver/receivers return either a first acknowledgement state in a second time slot, after the first time slot, or a second acknowledgement state in a third time slot after the second time slot.

One of the acknowledgment states is a positive acknowledge and the other acknowledgment state is a negative acknowledge.

Preferably, the first acknowledgment state is the positive acknowledge, and the second acknowledgment state is the negative acknowledge.

Preferably, the first time slot is variable in length and the second and third time slots are fixed in length.

Preferably, upon each transceiver/receiver detecting a correctly coded transmission in the negative acknowledge time slot, each transceiver/receiver discards the data previously received in the first time slot, and the transceiver/transmitter retransmits the data to each of the transceiver/receivers.

According to a third aspect of the present invention, there is provided a transceiver/receiver for use in a radio communications system including at least one transceiver/transmitter and at least one other transceiver/receiver, in use, the transceiver/receiver upon receiving a data packet in a first time slot from said transceiver/transmitter, either transmits a first acknowledgment state in a second time slot, after the first time slot, or transmits a second acknowledgment state in a third time slot, after said second time slot.

Preferably, the transceiver/receiver further receives the first acknowledgment state in the second time slot from the at least one other transceiver/receiver in the communication system or receives the second acknowledgment state in the third time slot from the at least one other transceiver/receiver in the communication system.

One of the acknowledgment states is a positive acknowledge and the other acknowledgment state is a negative acknowledge.

Preferably, the first acknowledgment state is the positive acknowledge, and the second acknowledgment state is the negative acknowledge.

Upon receiving a negative acknowledge from the at least one other transceiver/receiver, the transceiver/receiver discards the data packet received in the first time slot.

According to a fourth aspect of the present invention, there is provided a transceiver/transmitter for use in a communications system including at least one other transceiver/receiver, wherein in use, the transceiver/transmitter transmits a

data packet in a first time slot to the at least one transceiver/receivers in a first time slot and receives either a first acknowledge state in a second time slot, after the first time slot from one or more of the transceiver/receivers or receives a second acknowledgment state in a third time slot after the second time slot from at least one of the transceiver/receivers.

One of the acknowledgment states is a positive acknowledge and the other acknowledgment state is a negative acknowledge.

Preferably, upon receiving a negative acknowledge, the transceiver/transmitter retransmits the data to the at least one transceiver/receivers.

The system and protocol of the present invention has many uses including applications in controlling domestic, industrial and office appliances.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 shows one network architecture according to a preferred embodiment of the present invention;

Figure 2 shows the network protocol model used in the environment of the present invention;

Figure 3 shows a frame structure according to a preferred embodiment of the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

An exemplary architecture of a network is shown in Figure 1, in which the network 10 is made up from nodes 1, 2 and 3. Nodes 1, 2 and 3 are transceiving devices and may act as transmitters and/or receivers in a given communication Transaction. The network 10 may communicate with other networks 20, via gateway 11.

The protocol design of the present invention is based on the ISO seven layer model and some terminology is common with that used by ISO. The protocol used in the present invention is connectionless, meaning that once a single data transfer has

taken place, there is no expectation of additional related data transfers either before or after.

The protocol model of the present invention is based on the ISO seven layer model and is shown in Figure 2. For applications distributed over two nodes, each protocol layer has a virtual connection to the equivalent layer in the other node. As can be seen, each layer takes data provided by the layer above, treats it as a data unit and adds its own protocol control information (PCI) field. At each layer, the protocol data unit (PDU) is either the data, or a package provided by the next higher layer. The name of the PDU is prefixed by the layer to which it applies (For example, SPDU is a session PDU).

The physical layer relates to the mechanical and electrical network interface in the ISO system. In the system of the present invention, the physical layer refers to the hardware and firmware elements used to transmit and receive bits over the communication medium.

In the ISO system, the link layer is used for data link control (for example, framing, data transparency, error control). In the present invention, the link layer is used to break bytes up into bits, bit stuffing (if required), framing, collision detection, prioritisation, error detection, positive/negative acknowledge generation, checking and retransmission.

The network layer in the ISO system is used for networking routing, addressing, call set-up and clearing while in the present invention, the network layer is used for network routing, addressing, Transaction set-up and clear.

In the ISO system, the transport layer is used for end to end message transfer, connection management, error control, fragmentation and flow control. The transport layer is not used in the environment of the present invention.

The session layer in the ISO system is used for dialogue and synchronisation control for application entities but is not used in the environment of the present invention.

The presentation layer is used for transfer syntax negotiation, and data representation transformations in the ISO system while in the environment of the present invention, the presentation layer is used for optional encryption of application data.

The application layer in the ISO system is used for file transfer, access management, document and message interchange, job transfer and manipulation while in the environment of the present invention, the application layer supports sending and receiving application data.

Finally, the user application layer is used both in the ISO and the environment of the present invention for whatever is needed to achieve a specified function or behaviour.

It is predominantly in the link layer that the features of the present invention reside.

In the protocol of the present invention, use can optionally be made of a dominant bit, and an inferior bit. If two devices simultaneously transmit a dominant and an inferior bit, then receivers and transmitters (monitoring their own transmissions)will detect only the dominant bit.

Media access is obtained by a transmitter first monitoring the media, and if no existing transmission is detected, the transmitter will try to claim media access by transmitting a preamble stream. This preamble starts with at least one detectable bit. The claim for media access defines the start of a Transaction. A Transaction consists of all data transfer, acknowledgement and repeating of data. All nodes in a network must monitor the media continually and if they detect a Transaction occurring they will defer any attempt to claim media access until the completion of the current Transaction.

Transactions are asynchronous: they can occur at any time and the time difference from the start of one Transaction to the next does not have to be an integral number of bit periods.

In this application, a Transaction is specifically defined as a continuous period of time broken up into several sub-time slots containing different types of data. A Transaction will begin with a preamble for a set period of time, followed by the specific data which is to be transmitted from a transceiver/transmitter to two or more transceiver/receivers. The timeslot during which the data is transmitted is variable in length, and includes a portion used as a frame check sequence. Following the data transmission is a repeat timeslot tag (details of which will not be described as they are not relevant to the present invention) followed by a timeslot during which positive acknowledgement can be transmitted by the transceiver/receivers, followed by a timeslot during which negative acknowledgement is transmitted by the transceiver/receivers. The structure of this frame is shown in Figure 3.

As described above, a Transaction is asynchronous and can start at any time. However, once started, the Transaction has a time-based structure. Special markers in the Transaction are used to show the beginning and end of the data portion. The time slots during which positive and negative acknowledgement are transmitted, are fixed in time. By careful coding and redundancy of data encoded into these timeslots, a positive acknowledgement by one or more transceiver/receivers and a negative acknowledgement by one or more transceiver/receivers can be conveyed. All of the devices involved in the Transaction see both of the acknowledgment timeslots.

Transceiver/receivers wishing to positively acknowledge, will transmit a special code during the positive acknowledge timeslot and during the negative acknowledge timeslot will either receive (if dominant bit transmission is not used) or transmit inferior bits (if dominant / inferior bits are used).

Similarly, transceiver/receivers wishing to negatively acknowledge will either transmit inferior bits (if dominant / inferior bits are used) or receive (if dominant bit transmission is not used) during the positive acknowledge timeslot and transmit a special code during the negative acknowledge timeslot.

The fact that the devices monitor the timeslots they are not transmitting, ensures that by the end of two timeslots each device has detected either positive acknowledges, negative acknowledges or both and can therefore work-out the overall acknowledge state of the network.

For example, the transceiver/receiver which transmits a positive acknowledge will be able to detect some other transceiver/receiver which transmits a negative acknowledge. For the case where dominant bits are used, this is because the transceiver/receiver positive acknowledging will attempt to transmit inferior bits during the negative acknowledge timeslot but will detect dominant bits due to the other transceiver/receiver which is simultaneously transmitting an appropriate code during the negative acknowledge timeslot. For the case where dominant bits are not used, the transceiver/receiver positive acknowledging will receive during the negative acknowledge timeslot and detect any bits due to the other transceiver/receiver which is transmitting an appropriate code during the negative acknowledge timeslot. The reverse case applies for the transceiver/receiver transmitting a negative acknowledge.

At the end of the Transaction, all devices do not know how many positive or negative acknowledges there were, all they have to know is that there were some positive and some negative.

If there were any negative acknowledges at all during the Transaction, then all of the transceiver/receivers know this, and can discard the data received. Similarly, the transceiver/transmitter knows this and can attempt to re-run the Transaction.

The generation of a positive acknowledge will be as follows. Upon receiving data, a node will generate a positive acknowledge only when:

-the data timeslot has been checked against its embedded frame check sequence and found to be valid; and

-any addressing information present in the data timeslot matches an addressing information used by the device.

Each device transceiver generally contains at least two different types of address, as follows:

- A Unit Address, allowing the device to be uniquely addressed in isolation; and
- A Multicast address, allowing those devices in a network to be addressed simultaneously for the purpose of updating shared network variables.

In addition, devices can also optionally contain:

- A Network Address, allowing physical devices to be grouped by the logical network to which they are allocated.

Other variations are possible, but these three address types are fundamental, and used as the basis for other more sophisticated addressing schemes.

The processes involved in generation of a negative acknowledge are as follows. A receiving device (transceiver/receiver) will generate a negative acknowledge only if the data timeslot is determined to be corrupted, by checking the received data using the embedded frame check sequence.

Where a device determines data timeslot corruption, there is no point further examining any fields inside the data timeslot.

Acceptance of the data transmitted by the transceiver/transmitter is only made by the transceiver/receivers if the condition for generation of a positive acknowledge are met and if no other transceiver/receiver has generated a negative acknowledge. This ensures that all transceiver/receivers receive a given message only once. For point-multipoint messages, this may mean that a message is discarded by a transceiver/receiver, even if it appears valid and was positively acknowledged.

It will be understood that the above has been described with reference to a preferred embodiment and that many variations and modifications may be made within the scope of the present invention.

DATED this 8th day of August 2003.

Clipsal Integrated Systems Pty Ltd
By its Patent Attorneys
MADDERNS

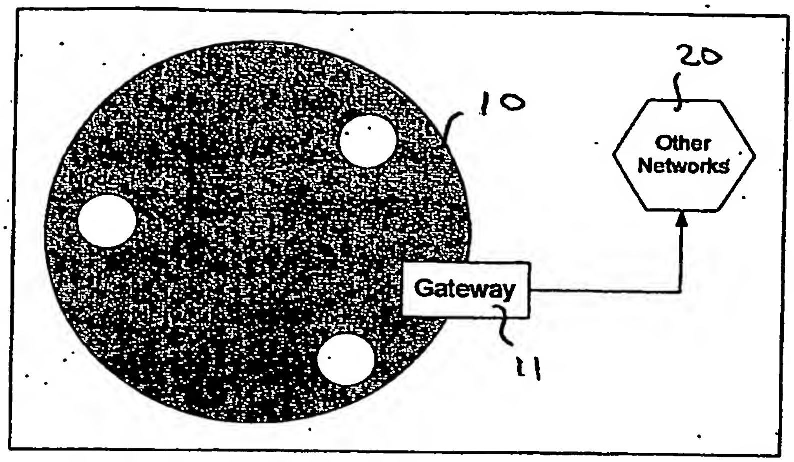

Fig. 1

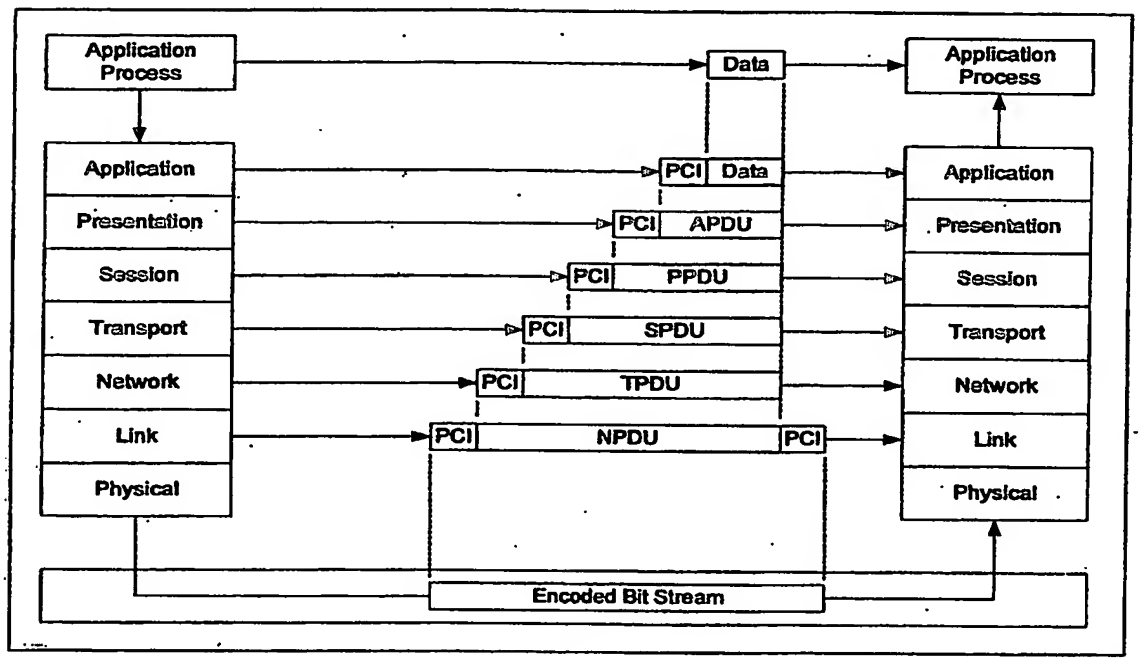

Fig. 2

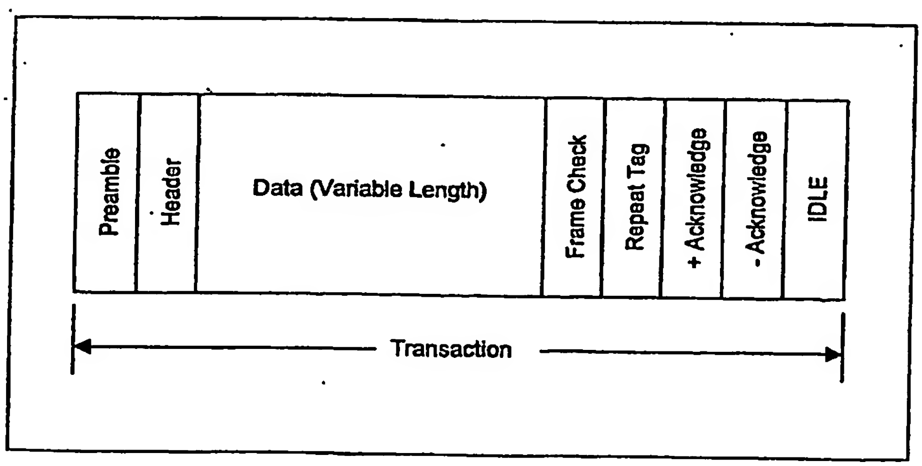

Fig. 3